B EDM Backup Configuration File

The backup configuration file (/usr/epoch/EB/config/eb.cfg) on the EDM Backup server defines how backups run at your site. The configuration interface stores its data in the eb.cfg file.

You configure your EDM backups using the Backup Configuration wizard, and you tailor your configurations to your needs by using the EDM Backup Configuration window. (Avoid manually editing the configuration file.) When you do not have access to the graphical interface (such as when dialing in), you can use the interactive command line interface, eb_config.

WARNING: Manually editing the eb.cfg file risks corrupting the backup configuration.

This chapter describes the statements within the configuration file. These statements correspond to windows and fields in the EDM Backup Configuration wizard and window, but they do not match exactly in wording or in number. This chapter does not decribe Symmetrix Path, Symmetrix Connect, or new database clients.

The topics in this chapter include:

- General Coding Rules
- Summary of Fields
- Server Fields
- Work Group Fields
- Filesystem Work Item Fields
- Database Work Item Fields
- PC Work Item Fields
- Backup Trailsets

General Coding Rules

If you must edit the configuration file, follow these general editing rules:

> **Note:** The EDM Backup Configuration window reformats the eb.cfg file. If you edit eb.cfg directly, you lose comments and spacing the next time you run the interface. The EDM Backup Configuration window reads in C pre-processor #include statements.

- The configuration file is made up of nested sets of blocks; a set of curly braces ({...}) delimits each block. Make sure to use these braces in pairs. For each open brace, there must be a closing brace.
- If you include multiple conflicting specifications for any given field, the last specification is used in most situations, unless otherwise noted in this chapter.
- Use a semicolon (;) to terminate entries that do not begin with a brace:

```
work group list: "dave's group","tony's group";
```

- Separate multiple strings by using a comma and optional spaces:

```
authorized backup list: "cad1", "cad2", "cad3",
"atlas2";
```

- Do not use the line-continuation character (\).

- Any comments are delimited using paired slash-asterisk characters: /**comments**/. Comments should not span lines.
- Place double quotes around *strings*, where a string is any non-numeric value except the reserved words that are described in the next bullet.

  ```
  client backup username: "ebadmin";
  ```

  Each quoted string must be complete on a single line: strings cannot span lines. Make sure to specify each string by using the correct case (capitalization). EDM Backup is case-sensitive when interpreting strings.
- Some reserved words do not require quotes; you can enter them as follows:
  - Always use the standard three-character abbreviation when specifying a month. Use all lowercase characters, or uppercase the first character only (for example, *mar* or *Mar*).
  - Specify units of time as *hours*, *minutes*, *seconds*, *days*, *weeks*, etc., as described for each field. Always use lowercase characters when specifying units of time.
  - Spell out days of the week completely. Use all lowercase characters, or uppercase the first character only (for example, *sunday* or *Sunday*).
  - Specify units of computer storage as follows: KB or K; MB or M; or GB or G (lowercase or uppercase).

Checking Your Changes

If you do edit the configuration file, you always must check afterward for syntax errors. To do this, run **ebbackup** with the **-check** option. This causes **ebbackup** to report any syntax errors.[1]

1. Under certain unusual conditions, some sytax errors will not be detected until the actual backup is run.

As necessary, edit the file to correct the errors, and then run the check again. Repeat this process until no more errors are present.

When Changes Take Effect

When you change various configuration parameters, changes either:

- take effect immediately on your system (even as backups are being processed), or
- take effect the next time **ebbackup** processing is started, either from an entry in root's crontab file or by issuing the **ebbackup** command manually

This does not include the next time **ebbackup** autoscheduling-only (**-sched** option) or checking (**-check** option) is started. Any exceptions are noted in the discussions.

Summary of Fields

Table B-1 summarizes the fields in the configuration file. It includes a brief description of each field, then tells whether the field is required (Req column), the field length and range of possible values, the initial value at installation, and when a change to the field takes effect.

The fields are presented generally in the order that they appear in the file.

Table B-1 Summary of Fields in the Configuration File

Field	Description	Req	Size or Values	Initial Setting	Takes Effect	
					Now	Next Bkup
Server Block Fields						
ebserver	Name of the EDM Backup server	yes	1-63 char	server-name	—	—
client backup username	Login name used when **ebbackup** connects with client system	yes	1-63 char	ebadmin		✔
backup administrator usernames	Individuals who are responsible for EDM Backup, and who have root-like privileges within the EDM Backup environment	yes	1-63 char (per user)	root	✔	
authorized backup list	Client systems that are backed up by the EDM Backup server	yes	1-63 char (per client)	none	✔	
authorized recovery list	Users who can restore files backed up from their own client (known as a *self-service restore*)	no	*client:user* combinations	none	next restore	
authorized cross recovery list	Users who can restore files to clients other than those from which the backup files originated	no	*client:user* combinations	none	next restore	
recovery administrator list	Users who can restore files other than their own to the client from which the files were originally backed up	no	*client:user* combinations	none	next restore	
maximum simultaneous clients	Global maximum number of clients that are backed up concurrently across all trails	yes	in the range 1-64	24	✔	

Table B-1 Summary of Fields in the Configuration File (Continued)

Field	Description	Req	Size or Values	Initial Setting	Takes Effect	
					Now	Next Bkup
use at most... trails concurrently	Maximum number of trails EDM Backup can write to a specific type of media concurrently	yes	*nn media_type*	4 per type of device	✔	
limit throughput	Total throughput of backup processing	yes	no limit *nnn K* per hour *nnn K* per minute *nnn K* per second (*K* is KB, MB, or GB)	600KB per second	✔	
maximum server backups.log file size	Maximum size of the server's backups.log file	yes	*nnn K or* no limit (*K* is KB, MB, or GB)	256KB	✔	
maximum server recoveries.log file size	Maximum size of the server's recoveries.log file	yes	*nnn K or* no limit (*K* is KB, MB, or GB)	256KB		next restore
maximum client backups.log file size	Maximum size of each client's backups.log file	yes	*nnn K or* no limit (*K* is KB, MB, or GB)	64KB		next backup for the client
maximum client recoveries.log file size	Maximum size of each client's recoveries.log file	yes	*nnn K or* no limit (*K* is KB, MB, or GB)	64KB		next backup for the client
catalog threshold to force level 0 backup	Level 0 backup is forced for any filesystem work item when more than this number of files in the work item have never been backed up	no	number of files; 0 disables this feature	0		✔
Work Group Field						
work group	Name of the work group	yes	1-63 char	none		✔
File System Work Item Fields						

Table B-1 Summary of Fields in the Configuration File (Continued)

Field	Description	Req	Size or Values	Initial Setting	Takes Effect	
					Now	Next Bkup
work item	Name of the work item, the client backed up by the work item	yes	1-63 char (each, per work item & client name)	none		✔
filespec	Name of one or more client directories, filesystems, or files that receives level 0-9 backups via this work item	yes	pathnames and optional **findxcpio** qualifiers (no size limit)	none		✔
baseline filespec	Same as above, but lists those local-client directories and filesystems in HSM systems that receive a baseline backup	no	same as *filespec*	none		✔
migration backup tag	Tag used to cross-reference a migration store on the server with the migration-client files from which the store was created (applicable when backing up staged files from the server)	no	1-63 char; first 16 must be unique	none		✔
exclusion tag	Marker that identifies work items that should not be backed up concurrently	no	1-63 char	none		✔
connection via	Specifies the use of an alternate network port for a work item in a multiple networked client.	no	alternate port name, must be a valid client name	none		✔
priority	Priority at which the work item should run	no	in the range -25 (run first) to 50	0		✔
do not load balance	Statement used to exclude the work item from load balancing (so no extra level-0 backups are taken)	no	as stated	statement omitted (load balance)		✔

Table B-1 Summary of Fields in the Configuration File (Continued)

Field	Description	Req	Size or Values	Initial Setting	Takes Effect	
					Now	Next Bkup
completeness	Code used to limit the number of files for which the data portion is written to the backup media (applicable for files under migration control, to avoid performing excessive backups)	no	all files resident files only files not backed up in migration store non-baselined files only	varies by client		✔
level map	Mapping between backup levels that normally occur (based on the template used) and the level you want for this work item	no	12-char string: example: "Bx0xxxxxxxx9"	"Bx0123456789"		✔
maximum files not backed up before forcing full backup	Level 0 backup is forced for the filesystem work item when more than this number of files in the work item have never been backed up	no	number of files; 0 disables this feature	0		✔
Database Work Item Fields These "Database Work Items" pertain to the "offline database backup" functionality supported prior to EDM 4.4.0 and are included here to support restores only. They do not apply to the database backup clients currently supported for backup.						
work item	Name of the work item, the client backed up by the work item	yes	1-63 char (each, per work item & client name)	none		✔
filespec	Name of one or more client directories, filesystems, or files that receives level 0-9 backups via this work item	yes	pathnames and optional **findxcpio** qualifiers (no size limit)	none		✔
database type	Type of database to be shut down	yes	oracle, informix, sybase	none		✔

Table B-1 Summary of Fields in the Configuration File (Continued)

Field	Description	Req	Size or Values	Initial Setting	Takes Effect Now	Takes Effect Next Bkup
database server	Name of the database server for the work item	yes	server name	none		✔
database name	Name of database(s) found during configuration	yes	database names	none		✔
type	Type of work item	yes	coordinated	coordinated		✔
use connection method	Defines the TCP/IP port on which the client is listening for a command from the EDM Backup server. Do not use for kicker work items	yes	socket@*port*	socket@0		✔
backup client initialization command	Script to be run on the client before shutdown	yes	See "Database Work Item Fields" on page B-50	none		✔
backup client cleanup command	Script to be run on the client after restart	yes	See "Database Work Item Fields" on page B-50	none		✔
backup client data buffer size	Data buffer sizes on client and servers	yes	0-32 MB	0 MB		✔
backup server data buffer size		yes	0-32 MB	0 MB		✔
recovery client data buffer size		yes	0-32 MB	0 MB		✔
recovery server data buffer size		yes	0-32 MB	0 MB		✔

PC Work Item Fields
These fields are used for all PC clients: NetWare, Windows NT, and OS/2 and for OpenVMS Clients. In several cases the word netware is embedded in the code, but the field is used for all PC clients.

Table B-1 **Summary of Fields in the Configuration File (Continued)**

Field	Description	Req	Size or Values	Initial Setting	Takes Effect	
					Now	Next Bkup
work item	Name of the work item, the client backed up by the work item	yes	1-63 char (each, per work item & client name)	none		✔
filespec	Name of one or more client directories, filesystems, or files that receives level 0-9 backups via this work item	yes	pathnames and optional **fndxcpio** qualifiers (no size limit)	none		✔
exclusion tag	ensures single-threaded processing for the client. Required and automatically generated for DOS.	no		none		✔
use connection method	Defines the TCP/IP port on which the client is listening for a command from the EDM Backup server	yes	netware@*port* (see "Connection Method" on page B-57)	netware@1776 for NetWare netware@1492 for OS/2, netware@3895 for Windows NT netware@3896 for OpenVMS		✔
netware username	Determines what user privilege can execute backups and restores	yes	1-63 chars (for NT: max 20 characters)	none		✔
netware encrypted password	Encrypted password for netware username	yes	not editable directly. Enter from EDM Backup Configuration window	none		✔
netware client TSA	Assigns a target service agent to the PC server	yes	*servername* .FileSystem (.local for OS2)	none		✔

Table B-1 Summary of Fields in the Configuration File (Continued)

Field	Description	Req	Size or Values	Initial Setting	Takes Effect	
					Now	Next Bkup
netware client target	Defines the PC client as a target in need of service	yes	client name	none		✔
Backup Trailset (Media Set) Fields						
backup trailset	Name of the trailset	yes	1-63 char	primary		✔
use trail	Name and type of media used for the trail, level of backups written to the trail, and the maximum number of work items that can be backed up concurrently to this trail	yes	name: 1-11 char type: *media_type* level: 0-9, B1, B2 max clients: 1-24	none media_type none 8	✔ max client	✔ other fields
use level *n* for baseline backups	Level of baseline backup written to this trail set (applicable with backups of HSM systems that are scheduled automatically)	no	B1 *or* B2	none	next time scheduling occurs	
keep backups	Retention period for the backup media, qualified by level (0-9)	no	*nn* months *or* forever (*months* is seconds, days, weeks, months, or years)	1 yr (lev 0) 3 mos (lev 1-9)	✔	
keep backup catalogs	Retention period for backup catalogs, qualified by level (0-9)	no	same as for *keep backups*	1 yr (lev 0) 3 mos (lev 1-9)	✔	
keep saveset records	Retention period for saveset records, qualified by level (0-9, B1, B2)	no	same as for *keep backups*	1 yr (lev 0) 3 mos (lev 1-9)	✔	
backup catalog delta level	Backup level at which EDM Backup should consolidate catalogs	no	1-9	9	✔	

Table B-1 Summary of Fields in the Configuration File (Continued)

Field	Description	Req	Size or Values	Initial Setting	Takes Effect Now	Takes Effect Next Bkup
duplicate media after backup on ..copies	Automatic media duplication	no	1	none		✔
manual activation of media duplication	Manual media duplication	no		none		✔
appending to current media copy	Append to current duplicate media for media duplication, as opposed to use new media.	no		none		✔
using new media at each duplication	Use new media for each duplication	no		none		✔
Backup Schedule Template Fields						
backup template	Name of the schedule template	yes	1-63 char	default		✔
work group list	Work groups that are backed up using this schedule template	yes	1-63 char (per work group)	none		✔
begin trailset rotations on	Date in future when EDM Backup should start using the trailset	yes	*dd-mmm-yy* *mm/dd/yy* *mmm dd, yy* *mmm dd, yyyy*	date template is added	next time scheduling occurs	
rotation period	Period of time during which each client should receive at least one level-0 backup	yes	*nn* days (*days* is days, weeks, months, or years)	14 days	next time scheduling occurs	

Table B-1 Summary of Fields in the Configuration File (Continued)

Field	Description	Req	Size or Values	Initial Setting	Takes Effect	
					Now	Next Bkup
primary trailset	Name of the trailset to which backups are written for all the work groups backed up using this template	yes	1-63 char	primary		✔
alternate trailset	Name of the trailset used to provide an optional, second set of backups on alternating nights (generally only used with automatic scheduling)	no	1-63 char	none		✔
logging level	Level of logging messages written to the file specified via the *server log file* field	yes	none errors stats debug per file	stats	✔	
server log file	Name and maximum size of the template's log file (stored on the server as /usr/epoch/EB/log/*filename*)	no	pathname & file size size: *mm* K or no limit (*K* is KB, MB, or GB)	default_ template.log 256KB	✔	
backup completion script	Name of the script file used to store or mail backup completion reports	no	pathname	mailok	✔	
backup failure script	Name of the script file used to store or mail backup failure reports	no	pathname	mailerr	✔	
do all baseline backups before normal backups	Statement used to force all scheduled baseline backups to finish before any level 0-9 backups start for the template	no	as stated	statement omitted (don't force baselines first)		✔

Table B-1 Summary of Fields in the Configuration File (Continued)

Field	Description	Req	Size or Values	Initial Setting	Takes Effect	
					Now	Next Bkup
recreate baseline if needed	Statement used to re-baseline files automatically, if the baseline copy is used in place of a staged-out copy during restore	no	as stated	statement omitted (do not recreate baseline)		✔
schedule standard rotations	Statement used to turn on automatic scheduling, and to schedule some portion of the clients for a full backup each day (cannot be specified with *full during weekends rotations*)	no	as stated	none	next time scheduling occurs	
schedule full during weekends rotations	Statement used to turn on automatic scheduling, and to schedule all full backups on Saturdays and Sundays (cannot be specified with *standard rotations*)	no	as stated	none	next time scheduling occurs	
weekday backup shift	Target amount of time EDM Backup should run each weekday (Monday-Friday; applicable with automatic scheduling to provide a guideline)	no	*hh* hours *mm* minutes in the range 1-24 hours	none		✔
weekend backup shift	Target amount of time EDM Backup should run each weekend day (Saturday-Sunday; applicable with automatic scheduling to provide a guideline)	no	*hh* hours *mm* minutes in the range 1-24 hours	none		✔

Table B-1 Summary of Fields in the Configuration File (Continued)

Field	Description	Req	Size or Values	Initial Setting	Takes Effect Now	Takes Effect Next Bkup
[for *work group*] level *n* on ...	Statement(s) used to turn on custom scheduling, and to specify when to run the backups for each level, optionally qualified by work group	no	work grp: 1-63 char level: 0-9, B1, *or* B2 on: days (*days* is a number [1 or greater] within the rotation, as in 1, 3-7, or the actual days, as in Monday, Tuesday, 2nd Monday)	none	next time scheduling occurs	
Startup Parameters						
perform initial full backup on scheduled day	Statement used to perform the initial full backup on some portion of all newly installed clients each day during the first rotation period (applicable with automatic scheduling)	no	as stated	statement enabled	next time scheduling occurs	
perform initial full backup as soon as possible	Statement used to run the initial full backup on all newly-installed clients during the first backup run after those clients are installed (applicable with automatic scheduling, and suggested if all new clients should receive a full backup as soon as possible)	no	as stated	statement enabled	next time scheduling occurs	

Server Fields

The server fields apply to all backup and restore operations that the server performs. The server block looks similar to the following:

```
ebserver: "atlas1"
{
client backup username: "ebadmin";
backup administrator usernames: "root", "jan",
"tracy";

authorized backup list: "cad1", "cad2", "cad3",
"cad4", "cad5","cad6", "doc1", "doc2",
"filserver1","atlas1";

authorized recovery list: "cad1:pat", "cad2:ken",
"cad3:eric", "cad4:jane", "cad5:tom", "cad6:bob",
"doc1:mary", "doc2:dave";

authorized cross recovery list: "cad1:pat<cad4",
"cad2:ken<cad6";

recovery administrator list: "*:jane", "cad2:ken";

maximum simultaneous clients: 4;
use at most 2 dlt trails concurrently;
limit throughput to: 600k per second;

maximum server backups.log file size: 256k;
maximum server recoveries.log file size: 256k;
maximum client backups.log file size: 64k;
maximum client recoveries.log file size: no limit;
catalog threshold to force level 0: 0;
.
.
.
}
```

The fields are described in order as shown above.

ebserver

This field specifies the name of the backup server, and defaults correctly when the server is configured at installation (with the **eb_server_config** command):

```
ebserver: "atlas1"
```

Client Backup Username

Use this field to specify the 1- to 64-character login name that **ebbackup** uses when it connects with client systems:

```
client backup username: "ebadmin";
```

Always use the installation default, ebadmin; this non-root user name is installed in /etc/password (or the NIS password map) for every client and server.

> CAUTION: **The ebadmin user name must have /bin/sh as its shell. Using any other shell causes installation and/or backup failures.**

If the default conflicts with another login name on your network, use an account that is dedicated to doing backups. It must be common to the server and all clients, and it cannot be a regular user name.

If your site is running more than one server, specify the same login name in the configuration file for each server.

> CAUTION: **If you change the client backup username, you must reinstall the client software on every client, including the local client on the server.**

Changes to this field take effect the next time **ebbackup** processing runs.

Backup Administrator Usernames

Use this field to identify individuals who are responsible for EDM Backup:

```
backup administrator usernames: "root", "jan",
"tracy";
```

You can specify any number of UNIX usernames (1 to 64 characters each, and set to "root" at installation). Each user whom you specify has root-like privileges *within the EDM Backup environment*.

The backup administrators can run backups and restore files from any client system, browse any backup catalog (regardless of restore permissions), perform cross-restores to restore any data to any client, and so on.

Note: Backup administrators do not necessarily need root access across the entire server system.

Repeat this field as many times as you want to add usernames. At a minimum, specify one name. Changes to this field take effect immediately.

Authorized Backup List

Use this list to identify those clients (workstations, fileservers, or backup or migration server) whose files can be backed up by this backup server:

```
authorized backup list: "cad1", "cad2", "cad3",
"cad4","cad5","cad6","doc1", "doc2",
"fileserver1","atlas1";
```

Each 1- to 63-character client name must match:

- the client name that is entered in the NIS map or the local /etc/hosts file
- the client name that is entered at installation. Make sure to use the expanded client name as registered in the NIS map (not an alias)

  and

- the client name that is specified in the work-item field

Repeat this field as many times as you want to add client names, or include all the backup clients in a single list. There is no default backup list. Unless you specify a client name here, it is not backed up.

Changes to this field take effect immediately.

Temporarily Disabling Backups for a Client

To disable backups temporarily for a client without changing its work-group definition(s), remove the client's name from this list. (You can still restore files backed up from the client.) When you put the client back in the list, backups resumes normally.

Authorized Recovery List

When you add a **Restore User Name** to the **Self Restore Permission** in the Backup Configuration window, that user name is added to the authorized recovery list in the eb.cfg file.

You can also add user names manually to the eb.cfg file. Use this field to identify users who can restore files that are backed up from their own client system, restoring them on their own (that is, the same) client:

```
authorized recovery list: "cad1:pat", "cad2:ken",
"cad3:eric","cad4:jane", "cad5:tom", "cad6:bob",
"doc1:mary", "doc2:dave";
```

This ensures that only authorized users can perform a self-service restore. Unless you specify this field, no users can restore their own files.

Repeat this field as many times as you want to add user names, or include all of the specifications in a single list. Backup administrators can always restore their own files, so they should not be specified explicitly in this list.

Identify each item in the list as a client:user pair, in this format:

```
authorized recovery list: "client:user",
..."client:user";
```

- Specify *client* as one of the following:
  - the name of the client system
  - an asterisk (*), which indicates any client

As an alternative to both the client and user names, specify the name of a netgroup preceded by an @, where the netgroup identifies each client:user combination that can perform a self-service restore:

```
authorized recovery list: "@doc";
```

- Specify *user* as one of the following:
  - the login name of the individual on that client who has permission to restore files to that client (illustrated below for cad1-cad3):

```
authorized recovery list: "cad1:*, "cad2:*",
"cad3:eric";
```

  - an asterisk (*), which means that any user who is logged into this client can restore the client's files:

```
authorized recovery list: "cad6:*";
```

    as an equivalent to the asterisk, omit the colon and username completely:

```
authorized recovery list: "cad6";
```

Changes to this field take effect the next time a restore runs.

You can also disable self-service restores for all clients by leaving the list blank (or by omitting the list):

```
authorized recovery list: ;
```

Authorized Cross-Recovery List

Use this field to identify users who can restore files to client systems other than those from which the backups originated:

```
authorized cross recovery list: "cad1:pat<cad4",
"cad2:ken<cad6";
```

Users who request a cross-restore must either own the files that are being restored, or must have *read* access to those files under UNIX. Users must also have write and execute permissions to the destination directory.

Repeat this field as many times as you want to add user names, or include all the specifications in a single list. Do not include backup administrators in this list. Backup administrators can always restore from any client system to any other client system, regardless of the system on which they are working.

Specify the list using the format shown below:

```
authorized cross recovery list: "client:user <
backup_client";
```

Note: The user requesting a cross-restore becomes the owner of the restored files, regardless of the client from which they were originally backed up. An exception occurs if the user requesting the cross-restore is also a restore administrator, in which case the original owner is preserved.

- Specify *client* as the system to which files can be restored through a cross-restore. At restore time, this must be the client at which the person requests the cross-restore is logged in. Specify:
  - the name of the client system to which files can be restored
  - an asterisk (*), which indicates any backup client (illustrated below) where Pat can restore any files to (and from) any client system):

```
authorized cross recovery list: "*:pat < *";
```

  As an alternative to both the client and user names, specify the name of a netgroup preceded by an @, where the netgroup identifies each user that has permission to restore files to the clients in the netgroup:

```
authorized cross recovery list:"@netgroupa< cad5";
```

- Specify *user* as the name of the user authorized to initiate restore on the client identified to the left of the colon. You can specify:

  - the login name of the individual who has permission to restore files to that client, as it appears in the authorized recovery list. The following example gives the user Pat permission to restore files that are backed up from cad6 to cad1:

```
authorized cross recovery list: "cad1:pat < cad6";
```

  - an asterisk (*), which means that any user who is logged into the client can perform the cross-restore:

```
authorized cross recovery list: "cad1:* < cad6";
```

    as an equivalent to the asterisk, omit the colon and username completely:

```
authorized cross recovery list: "cad1 < cad6";
```

- Specify *backup_client* as one of the following:
  - the name of a client from which files were originally backed up (where the files can be restored to the indicated client:user)
  - an asterisk (*), which indicates any backup client
  - the name of a netgroup that is preceded by an @, to specify the clients in an entire netgroup

The following example allows any user who is logged in to cad1 to restore files from any backup client:

```
authorized cross recovery list: "cad1 < *";
```

Changes to this field take effect the next time a restore runs.

Disabling All Cross-Restores

Select **Restrict Cross-Restore to Clients** in the Client tab in the EDM Backup Configuration window to disable all cross-restores, except those that the restore administrator list explicitly allows, and those that the backup administrator performs.

To do this manually in the eb.cfg file, leave this list blank (or omit the list) to disable all cross-restores:

```
authorized cross recovery list: ;
```

Recovery Administrator List

Use this field to identify users who can restore files other than their own to the client from which the files were originally backed up:

```
recovery administrator list: "*:jane", "csd2:ken";
```

The user who requests the restore must be logged in to the client from which the files were backed up.

This feature is designed for use when multiple users share a client, to enable a single user on that client to serve as the restore administrator for just that client. Repeat this field as many times as you want to add usernames, or include all the specifications in a single list. Backup administrators have all of the rights of a restore administrator, so should not be specified explicitly in this list.

Identify each item in the list as a client:user pair, in this format:

```
recovery administrator list: "client:user",
..."client:user";
```

Note: The original owner is preserved for all files that a restore administrator restores.

- Specify *client* as one of the following:
  - the name of the client. This is the client from which the files were backed up and to which they can be restored by the restore administrator.
  - an asterisk (*) to indicate any client:

```
recovery administrator list: "*:karen";
```

  - the name of a netgroup preceded by an @, where the netgroup identifies each client in the netgroup.
- Specify *user* as the name of the restore administrator authorized to initiate restore on the client(s) indicated. You can specify:
  - the login name of the user.
  - an asterisk (*) to indicate any user:

```
recovery administrator list: "csd1:*";
```

As an equivalent to the asterisk, omit the colon and username completely:

```
recovery administrator list: "cadl";
```

Changes to this field take effect the next time a restore runs.

Disabling Administrator-Level Restores

Leave this list blank (or omit the list) to disable all administrator-level restore functions except those that the authorized cross-recovery list explicitly allows, and those that the backup administrator performs:

```
recovery administrator list: ;
```

Maximum Simultaneous Clients

Use this field to specify the *global* maximum number of work items (not clients) the EDM Backup server can back up concurrently, *across all trails* being written at any one time. This field controls the total amount of resources the server can allocate to backup functions, and relates to the server's available processing power, memory, and connectivity. Adjust this field to reduce or expand the system resources available for backup processing:

```
maximum simultaneous clients: 24;
```

Specify a value in the range 1-64 (set to 24 at installation). If you specify a value that is too small, EDM Backup under-utilizes system resources. If you specify a value that is too large, EDM Backup saturates the server's virtual memory, which causes memory thrashing and reduces performance.

Changes to this field take effect immediately.

Use At Most *n media-type* Trails Concurrently

Your server only has a fixed number of physical drives for backup. Besides EDM Backup, other applications on the server may need to use these drives, such as the HSM server.

Specify one *use at most* field for each type of backup media, to define the maximum number of drives EDM Backup can use for each device type (that is, the maximum number of trails EDM Backup can write to that type of media at once). This is initially set to 4 for each type of device that is available to the server.

This field has the format:

```
use at most n media-type trails concurrently;
```

where:

- *n* is an integer representing the number of server drives EDM Backup can use at any one time.
- *media-type* indicates the type of drive.

For example, if the EDM Backup server has four tape drives, but wants to reserve two drives for non-backup purposes, you specify:

```
use at most 2 dlt trails concurrently;
```

Changes to this field take effect immediately, although no backups are terminated to comply with the new limits.

Limit Throughput To: *nnn* Per *time*

Use this field to limit the amount of the network bandwidth available to EDM Backup, thereby allowing the network to accommodate applications other than EDM Backup. With this field specified, EDM Backup monitors its network use. If it reaches the specified limit, it does not start another client backup until the throughput drops. (However, EDM Backup does not stop any backups currently in process if the throughput exceeds this limit.) Specify *no limit* if you don't want EDM Backup to monitor the network.

Specify this field in terms of the number of bytes EDM Backup can use during a specific time period:

```
limit throughput to: bytes per time-units;
```

Where:

- *bytes* is the number of bytes to which you want to limit the throughput, followed by a unit code (set to 600KB per second at installation). Specify the unit code as follows:

Unit Code	Measure
k, K, kb, or KB	kilobytes
m, M, mb, or MB	megabytes
g, G, gb, or GB	gigabytes

- *time-units* defines the unit of time during which you want to limit throughput to the number of bytes specified:

Time Code	Limits Throughput Based on the Specified Number of Bytes Per:
second(s)	second
minute(s)	minute
hour(s)	hour

You can combine the time units, as in this example:

```
limit throughput to: 15MB per 1 hour 30
minutes;
```

Here are some guidelines:

- If the server is connected to a single Ethernet, set this to 6 (600KB/second).
- If you have high end clients, two independent FDDI rings, 6 CPU SC-1000, and 8 DLT drives, you can set this for up to 200 (20MB/second).

Changes to this field take effect immediately.

Specifying No Throughput Limit

If you want backups to proceed as quickly as possible with no restrictions on throughput, specify *no limit* (and no other options). You may want to do this if you have FDDI or multiple Ethernets.

```
limit throughput to: no limit;
```

Maximum Server backups.log File Size

Use this field to specify the maximum size of the backup log file (/usr/epoch/EB/log/backups.log) on the server (generally in the range 16KB-256KB). This file provides a record of all backup activities, limited only according to the size restriction that is specified here. Determining this size is a trade-off between using disk space on the server versus keeping the backup history for a longer period of time.

When the file reaches the specified size, EDM Backup locates the oldest data in the file and expires ten percent of that data to free up space for new information.

The format of this field is:

```
maximum server backups.log file size: file-size;
```

Where *file-size* indicates how large the file can get before the oldest data is expired. Specify this field as a number of bytes followed by a unit code (set to 256KB at installation). Specify the unit code as follows:

Unit Code	Measure
k, K, kb, or KB	Kilobytes
m, M, mb, or MB	Megabytes
g, G, gb, or GB	Gigabytes

If you want to let the file grow until it is as large as the physical storage space allows, specify *no limit*. By not setting a limit on the file, it becomes a permanent record of backup operations.

```
maximum server backups.log file size: no limit;
```

This example sets the maximum file size for backups.log to 200KB:

```
maximum server backups.log file size: 200KB;
```

Note: If you do not limit the file size, it may become too large to manage. Also, in HSM systems it is not staged.

Changes to size of the backup log file take effect immediately.

Maximum Server recoveries.log File Size

Use this field to specify the maximum size of the restore log file (/usr/epoch/EB/log/recoveries.log) on the server (set to 256KB at installation). This file provides an audit trail of all restore activities, limited only according to the size restriction specified here. It can be used to examine exactly what occurred during restore processing. Determining this size is a trade-off between using disk space on the server versus keeping audit trails on the server for a longer period of time.

The format of this field is:

```
maximum server recoveries.log file size: file-size;
```

For example:

```
maximum server recoveries.log file size: 200KB;
```

Specify this field exactly as described for the server's backup log file. With the exception of when changes take effect, these two settings are handled exactly the same.

Changes to the size of the restore log file take effect the next time a restore runs.

Maximum Client backups.log File Size

Use this field to specify the maximum size of the backup log file (~ebadmin/*client-name*/backups.log) on each client (set to 64KB at installation). This file provides a brief record of backup activities for the client. The format of this field is as follows.

```
maximum client backups.log file size: file-size;
```

For example:

```
maximum client backups.log file size: 64KB;
```

Specify this field exactly as is described for the server's backup log file. With the exception of the default and when changes take effect, these two settings are handled exactly the same.

Changes to this field take effect the next time the corresponding client is backed up.

Maximum Client recoveries.log File Size

Use this field to specify the maximum size of the restore log file (~ebadmin/*client-name*/recoveries.log) on each client (set to 64KB at installation). This file provides a brief record of restore activities for the client, and can be used to locate the more comprehensive restore information on the server.

The format of this field is:

```
maximum client recoveries.log file size: file-size;
```

For example:

```
maximum client recoveries.log file size: 32KB;
```

Specify this field exactly as is described for the server's backup log file. With the exception of the default and when changes take effect, these two settings are handled exactly the same.

Changes to this field take effect the next time the corresponding client is backed up.

Catalog Threshold to Force Level 0 Backup

This statement directs catalog processing to schedule a full (level 0) backup (instead of an incremental) for any filesystem work item when it detects too many files within that work item that have never been backed up. (Applies to filesystem backups only, as database backups are always full.)

The concern is for files that do not get backed up during an incremental backup because the files were added in a manner that preserved the creation and access time of the file prior to the last backup of the work item. In this case, you want the next backup for that work item to be a level 0. Failure to do this causes catalog processing to take a long time and could result in a large number of files not being backed up, meaning they could not be recovered.

Also, see "When You Change a Work Item or a Filesystem" on page B-35 for times you should force a level 0 backup.

The format of the field is as follows:

```
catalog threshold to force level 0 backup:
threshold;
```

For example:

```
catalog threshold to force level 0 backup: 10;
```

If **ebcatcomp** detects more than *threshold* files that have never been backed up for a filesystem work item, it schedules (command line scheduling) a level 0 for that work item. If *threshold* is set to zero, the default value, this feature is disabled.

Note: This statement can increase the number of level 0 backups performed.

The following alternate wording for this statement is acceptable in the server block:

```
maximum files not backed up before forcing full
backup: value;
```

This statement can be overridden by a version of this directive that can be set within individual filesystem work items. See "Maximum Files not Backed Up Before Forcing Full Backup" on page B-49.

Work Group Fields

Work groups define a set of work items of the same type so that a template can back them up together (during the same shift and using the same trailset) just by referencing their work-group name. Specify as many work groups as necessary to configure your site. If all your clients can be backed up together by using a single trailset, specify one work group that includes all client systems.

The work-group block looks similar to the following:

```
work group: "doc"
{
work-item definitions ...
}
```

Note: A work group can be referenced by more than one template. While this may result in backing up the clients in the work group more often than necessary, it does not cause any conflicts (in media, scheduling, etc.).

Work items in a work group must be of the same type. For example, a work group cannot contain file system work items with NetWare work items or database work items.

Specify the work-group name as described below, then proceed to the discussion on coding work items.

Work Group Name

Specify the work-group field as any unique 1-63 character name used to reference a group of clients, in the following format:

```
work group: "work-group name"
```

Note: The work-group name field does not end with a semicolon, but is followed by a brace-delimited block that defines its work items.

Changes to the work-group name take effect the next time **ebbackup** processing runs.

Filesystem Work Item Fields

Each work group is comprised of one or more *work items*. A file system work item defines a set of files to be backed up for a single client, and information relating to those files (whether to back up all the files, resident files only, and so forth). Each client requires at least one work item before it can be backed up by EDM Backup.

The following sample shows a work group. (The migration, baseline, and completeness fields are used with HSM.) The *local* work group backs up files on the EDM Backup server (named atlas1), and includes three work items:

- one work item backs up key database files

- two work items back up the non-database files from each of two filesystems (/ and /home).

Work item used to back up server database files

Work item used to back up the server's root (/) filesystem

Work item used to back up the server's /home filesystem

```
work group: "local"
{
  work item: "atlas1:LOCAL_DATABASE", "atlas1"
  {
   filespec: LOCAL_DATABASE;
   completeness:  DB_COMPLETENESS;
   priority: PRIORITY_SERVERDB;
   level map: LEVEL_MAP_SERVERDB;
  }
  work item: "atlas1:/", "atlas1"
  {
   filespec: "/ -xdev";
   exclusion tag: "/dev/rsd2C/c0t3d0s4";
   baseline filespec: "/ -xdev -staged";
   completeness: non-baselined files only;
  }
  work item: "atlas1:/home", "atlas1"
  {
   filespec: "/home -xdev";
   migration backup tag: "local_home"
   exclusion tag: "/dev/rsd2C/c0t3d0s4";
   baseline filespec: "/home -xdev -staged";
   do not load balance;
   completeness: non-baselined files only;
  }
}
```

Some of these options only apply for a particular type of EDM Backup client, as detailed in Table B-2.

Table B-2 **Work Item Options Available by Client Type**

Option	Applicable if the work item is:			
	on a Migration Client	on a Non-Migration Client	Local to the Backup Server	
			Level 0-9	Level B1 & B2
work item:	yes	yes	yes	yes
filespec:	yes	yes	yes	—
baseline filespec:	—	—	—	yes
migration backup tag:	—	—	yes[1]	—
exclusion tag:	yes	yes	yes	yes
priority:	yes	yes	yes	yes
do not load balance:	yes	yes	yes	—
completeness:				
all files	—[2]	yes[3]	yes	—
resident files only	yes[3]	—	yes	—
files not backed up in migration store	yes[4]	—	—	—
non-baselined files only	—	—	yes[4]	—
level map:	yes	yes	yes	yes

1. Applies when backing up staged files on an EDM Backup server that is also an EDM Migration server.
2. Only used when you reach a project milestone or temporarily decommission EDM Backup, to record the filesystems exactly as they stand at that time.
3. This setting can leave you vulnerable unless there's a backup of the client store. Specifically, if you don't back up a file that's been staged out, but you lose the file's client store on the server before the server's files are backed up, the only way you'll be able to restore the file is from an old backup.
4. Default for this client type.

Specify each work-item field as described in the following discussions. The fields are covered in order as shown in the syntax above. Repeat the *work-item* specifications as many times as necessary to configure backups for each client.

All changes that are made to the work item take effect the next time **ebbackup** processing runs.

When You Change a Work Item or a Filesystem

When you change a work item's file specification significantly (for example, to add a new filesystem to the backup list), always schedule the next backup for that work item as level 0, and make sure it is backed up the next time EDM Backup runs. Failure to do this causes catalog processing to take a long time and may result in a large number of files not being backed up.

Other times to force a level 0 are:

- When a filesystem covered by a work item changes its device number (such as when reformatted or when moved to a new system with the same hostname).
- When you have added files to a filesystem covered by a work item in a manner that preserved the creation and access time of the file prior to the last backup of the work item. See "Catalog Threshold to Force Level 0 Backup" on page B-29 and "Maximum Files not Backed Up Before Forcing Full Backup" on page B-49.

When You Stop Using a Work Item

If you stop using a work item — for example, after you create two work items for a client that previously had one — make sure to keep the old work item in the configuration file. If you delete that older work item, you lose its association with its client system, resulting in problems when you run history reports.

To avoid this situation, define one work group that contains only obsolete work items, and that is not backed up by any template. This retains the association between each old work item and its client, for each work item in the group.

Work Item Name

Specify a 1-63 character name that is unique within the configuration file. As a conventions, it is good to incorporate the name of the client when specifying the work-item name, as well as some indication of what files are backed up using the work item. EDM Backup generates work-item names that look similar to the following when each client is installed (shown for a work item that backs up *all* files on the *doc1* client):

```
work item: "doc1-all", "doc1"
```

Work item Name

At a minimum, there is one work item per disk; there can be several work items per client.

Note: In a HSM system, when defining local-client work items, always specify one work item to correspond to each client store that migrated to the server. Use that work item to back up the client store on the server. Make sure to include a Backup/HSM tag in the work-item definition, and reference the same tag in the corresponding store-name definition.

A work item can back up an unlimited number of files. However, a work item that backs up more than 250,000 files will have substantially slower catalog processing.

This example shows the doc1 client broken down into three work items:

```
work item: "doc1o/", "doc1"
{
   filespec: " ... ";
}
work item: "doc1o/usr", "doc1"
{
   filespec: " ... ";
}
work item: "doc1o/other", "doc1"
{
   filespec: " ... ";
}
```

Client Name

Specify the name of the client to which the work item applies, as it appears in the authorized backup list. Always use the primary complete name for the client; not an alias.

```
work item: "dool-all",  "dool"
```

Client Name

Filespec to Back Up

This field lists each directory, filesystem, and/or file that you want to back up using the work item. This field has a 4096 character limit.

Specify the exact names as they appear in a directory listing by using a syntax similar to **find**. For each directory or filesystem, include the name and optional qualifiers. Refer to Appendix D "findxcpio Directives" for more details about the syntax and the semantics of specifications that this field supports. (Also see the **findxcpio** man page.)

Using the Block Form of the Syntax (Filespec Statement)

The work-item syntax uses a separate statement (not on the same line as the work-item name).

```
work item:  "work-item name", "client name"
{
  filespec: "filespec to back up for level 0-9 backups";
  baseline filespec: "filespec to back up when baselining";
  migration backup tag: "migration backup template name";
  exclusion tag: "tag";
  priority: priority;
  do not load balance;
  completeness: completeness-code;
  level map: "level map";
}
```

Backing Up the Local Client

It is important that you back up all filesystems that are stored on your backup server. The server stores the database files for all EMC backup and HSM products, as well as the bitfiles (client stores) written out by EDM Migration. Also, the product installations modify several standard UNIX files, such as /etc/passwd, /etc/group, and crontab.

There is a special #define macro for use with the local client, which causes it to back up all the server's critical database files:

```
filespec: LOCAL_DATABASE;
```

The following macro is included in the autoconfigured work item for the server's database files (shown below for the atlas1 server), and should not be changed:

```
 work item: "atlas1:LOCAL_DATABASE", "atlas1"
{
   filespec: LOCAL_DATABASE;
   completeness:  DB_COMPLETENESS;
   priority: PRIORITY_SERVERDB;
   level map: LEVEL_MAP_SERVERDB;
}
```

Baseline Filespec

The baseline filespec is available if you purchased HSM, and applies when you back up the EDM Backup server's own files — generally only the files that are used for stageable filesystems. It lists each filesystem for which you want to maintain a baseline backup (level B1, B2, or both), in this format:

```
baseline filespec: "filespec to back up when baselining";
```

If you want to maintain a baseline backup for /home, you specify:

```
baseline filespec: "/home -xdev -staged";
```

Migration Backup Tag

Include a migration backup tag for use with level 0-9 backups on a local client, when:

1. The EDM Backup server is also an EDM Migration server and
2. The files that are being backed up are the staged (migrated) versions of client files

A migration backup tag is required in this case. Specify one work item for each client store, and include this field:
`migration backup tag: "migration backup template name";`

Where *migration backup template name* is a 1-63 character identifier. The first 16 characters must be unique across all work items that back up the server's own files, and must match *exactly* the migration *backup template name* that is specified in the definition of the migration store used to stage the same files from their network client.

Figure B-1 **Migration Backup Tag**

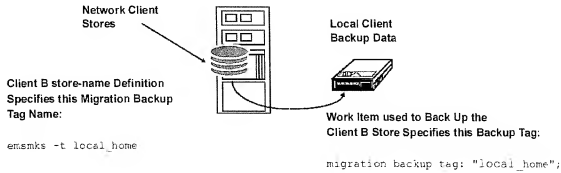

The migration backup tag is used to update the backupdates database (/usr/epoch/etc/backupdates), so that a subsequent backup of a network client that is migrated to this server can determine whether the staged-out client files were already backed up on their migration server (that is, the server where you define the work item).

Exclusion Tag

Use this field to identify any work item that should not be backed up concurrently with one or more other work items. For each client, specify the same tag value (1-63 characters) for all those work items that are mutually exclusive, using this format:

```
exclusion tag: "tag";
```

The combination of the exclusion tag and client name forms a unique identifier within the configuration file, but you can use the same tag value for any number of clients.

Typically, you use an exclusion tag for work items that back up data stored on the same physical disk, to prevent disk thrashing during backup processing (shown below for two work items that back up files from the rsd2C/c0t3d0s4 device):

```
work item: "atlas1:/", "atlas1",
 {
   filespec: "/ -xdev";
   exclusion tag: "/dev/rsd2C/c0t3d0s4";
.
.
work item: "atlas1:/home", "atlas1",
 {
   filespec: "/home -xdev";
   exclusion tag: "/dev/rsd2C/c0t3d0s4";
.
.
```

If no exclusion tag is included for a work item, **ebbackup** assumes it can process that work item concurrently with any other work item.

Connection Via

You can install just one copy of the EDM client software on UNIX client machines that are equipped with more than one Ethernet, FDDI, Token-Ring, or ATM port and use it to have your backups obtain higher throughput rates by taking advantage of the multiple network ports.

Setting up the Network

If possible, at least one of the network cards for the EDM and the clients should be connected to a separate network dedicated to backups. This reduces the amount of network traffic on the customers' main network and reduces the impact of backups on the users.

The EDM should be configured on the network in a way that each card uses a single/different client port. That is, card "A" on the server should talk to "A" on the client and card "B" on the server should talk to "B" on the client. Each client port must have a distinct clientname. For example: clientname_A.

Configuring the Work Items

You can configure the client's backups to be divided into separate backup streams (each defined by a "work item"). You can specify (in the work item) an alternative port name for one or more of the backup streams. The clientname is used as the name of the default port for backups. (It can match or differ from the hostname of the client machine.)

To set up a filesystem work item with an alternate network port, add the distinct clientname to the EDM Backup Configuration window: select **Work Item Options** and enter it in the **Generic** tab. To set up online database backups, use the **Configure Online Options** in the Client tab.

This puts the clientname in a **connection via** parameter in the work item statement in eb.cfg, which specifies it as the alternate port name.

Priority

Use this field to assign a priority to the work item, forcing it to run either before or after one or more other work items.

When you start a backup, **ebbackup** searches through its schedule for the work items that have the highest (that is, lowest-numbered) priority. It makes that priority *current*, and processes all work items that have that current priority to

completion. Then it searches for the next lower priority and makes it current, and so forth, until all work items are processed.

> **Note:** The use of priorities limits the flexibility that is available to EDM Backup in choosing what backups to run next, and might reduce the efficiency with which the server's resources are used.

If a work item is added to the schedule and has a priority that's higher (i.e., numerically lower) than the one being processed, that work item is started as soon as server resources are available, and its priority becomes the current priority. This might be the case, for example, if a request is submitted online, or **cron** starts a scheduled backup while another backup is still running.

Use this syntax to specify the priority:
`priority: priority;`

Where *priority* takes one of two forms:

- assigned integer in the range -25 to 50, where -25 specifies the highest priority (run first) and 50 specifies the lowest priority (run last).
- one of the following =define macros:

Table B-3 Priority Settings

Macro	Equivalent	Description
PRIORITY_FIRST	-25	Backups should run before the general backups.
PRIORITY_DEFAULT	0	Backups should be considered part of the general backups (no special priority). This is the initial setting at installation, except for the work item used to back up the server's database files.
PRIORITY_LAST	25	Backups should run after the general backups.
PRIORITY_SERVERDB	50	Backups should run last. This keyword is reserved for use in backing up the server's database files (LOCAL_DATABASE).

If the backup server is also an HSM server, EDM Backup automatically backs up the stores for each migration client before backing up the client files. You do not have to be concerned with setting a higher priority for the local-client work items in this case. That way, if the level 0-9 backup of a networked migration client calls for a backup of a migrated file, and the work item specifies a completeness option of *files not backed up in migration store*, the file do not have to be copied back in and backed up again.

To ensure that the migration stores are backed up first for each migration client, you must:

1. Use the same priority for the network client and all work items that back up the client stores on the server (generally priority 0, the default).
2. Set the migration tags correctly for all work items that back up the client stores on the server.

Note: If you purchased HSM (with baselining) and the template specifies *do all baseline backups before normal backups*, the baseline backups always run before any level 0-9 backups for that template, regardless of the priority for any specific work item.

Do Not Load Balance

Load balancing works in conjunction with automatic scheduling, and is the method whereby EDM Backup schedules backups as evenly as possible across all clients. By default, EDM Backup assumes that each work item participates in load balancing, which can result in some work items having more than one level 0 backup during any given rotation.

Include the following statement if you do not want the work item to participate in load balancing:

```
do not load balance;
```

This ensures that the work item does not receive more level 0 backups during the rotation period than are called for. This might be useful, for example, for filesystems that take a long time to back up, or for production systems whose resource consumption is critical.

Completeness

Use this option for files under migration control, to avoid performing duplicate backups of the same file data:

Note: You can use a special #define completeness macro, DB_COMPLETENESS, with the local client. This macro equates to *all files*. It is included in the autoconfigured work item for Epoch's database files and should not be changed.

Use this option for files under migration control, to avoid performing unnecessary backup. It limits the files for which the *data portion* is written to the backup. (Remember that the *extended inode* is included for each file scanned, regardless of whether its data is written out.)

```
completeness: completeness-code;
```

Specify the *completeness-code* as any of the values below. This field applies for level 0-9 backups. The initial setting varies depending on the type of file that is being backed up (as noted in Table B-4). Generally it is a good idea to use the defaults.

Table B-4 Completeness Settings in HSM Systems

Code	Description	Applicable For	Default For
All files	Back up the data portion of all files in the filespec, regardless of where they're stored and whether they're baselined.	All clients (this is the only option available for backup clients that are not also EDM Migration clients)	Non-migration clients & server's database files on the server
Resident files only [1]	Back up the data portion for only those files that are resident (local to) the client; or, for the server, that are stored on the magnetic disk.	Levels 0-9 on the backup server (i.e., the local client), and EDM Migration clients	—
Files not backed up in migration store	If a file has been staged, only back up the data portion of the file if its staged version hasn't been backed up yet on the HSM server.	EDM Migration clients	EDM Migration clients
Non-baselined files only	Back up the data portion for only those files that aren't baselined (for use after a baseline is taken). This option is only available if you've purchased HSM.	Local (server) client (but not server's database files)	Local client (except server's database files)

1. It's safest to use the next option (files not backed up in migration store) with EDM Migration network clients (and non-baselined files only for the local EDM Migration client). The resident files only setting can leave you vulnerable unless there's a backup of the client store. Specifically, if you don't back up a file that's been staged out, but you lose the file's client store on the server before the server's files are backed up, the only way you'll be able to restore the file is from an old backup.

Level Map

Use this option to override the backup level(s) that normally occur for the work item, based on the template that is used. This feature is useful when a few work items need a level of backup that is different from most of the work items referenced by a template, and saves having to define a separate template for those work items.

The level that a level map specifies completely replaces (in the schedule) the level that is otherwise performed. EDM Backup does not keep a record of the old (replaced) level.

Note: The level map applies for custom and automatic scheduling. Command-line scheduling does not use the level map.

Specify the level map by using one of the #define macros that are listed in the following table. For example:

```
level map: LEVEL_MAP_SKIP_ALL_BUT_L0;
```

Table B-5 Level Map Settings

Level Map #define Macro	Equivalent Value	Description
LEVEL_MAP_SKIP_ALL_BUT_L0	"BB0xxxxxxxxx"	Only run level-0 and baseline backups. Skip all other levels.
LEVEL_MAP_EVERYTHING_IS_L0	"BB0000000000"	Run all scheduled backups as level-0, regardless of the level specified by the template.
LEVEL_MAP_SKIP_ALL	"xxxxxxxxxxx"	Skip the work item completely. Don't run any backups.
LEVEL_MAP_SERVERDB [1]	"BB0000000000"	Same as LEVEL_MAP_EVERYTHING_IS_L0, but reserved for when backing up the EDM Backup server's database files.

1. The LOCAL-DATABASE work item (for EDM Backup database files) must always receive a level 0 backup, using the LEVEL_MAP_SERVERDB level map. This simplifies restoring those critical files.

Alternatively, specify any 12-character string you want, where each character is indexed to a backup level (B1, B2, "BB" then 0-9, in that order). Use an *x* to suppress a level of backup

completely. Any other value refers to the level of backup that should occur when the work item is processed at the indexed level.

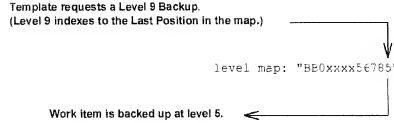

The following string requests that each backup proceed at the normally scheduled level, and indicates that there is only one baseline backup:

```
level map: "Bx0123456789";
```

Though you cannot remap a baseline backup to a different level, you can cause it to be ignored. If the trailset requests a B1 backup, for example, you cannot change it from the first id (level map: "Bx01...") to the second id (level map: "xB01..."), but you can suppress it entirely:

```
level map: "xx0123456789";
```

Here are some examples of level mapping:

- To change level 6-9 backups to a level 5 for the work item:
  ```
  level map: "BB0123455555";
  ```
- For stable filesystems that seldom change, run the baseline and level-0 backups as scheduled, but ignore other levels:
  ```
  level map: "BB0xxxxxxxxx";
  ```
- For filesystems that change frequently but for which incremental backups are not used during restore (as for the EDM Backup databases), run all backups as level 0:
  ```
  level map: "BB0000000000";
  ```

- If you are running level 9 backups manually during the day (in addition to the automatically scheduled nightly backups), you might map the nightly level 9 backups to level 5, then specify a separate trail (set of media) for the levels 5 and 9. The trail that is used to restore files do not involve the manually run level 9 incrementals that are taken throughout the day:
  ```
  level map: "BB0xxxx5xxx5";
  ```
  Because the level map is not used when you run command-line backups, the level 9 incrementals that are run from the command line actually perform level 9 backups.
- If you want to suppress all backups:
  ```
  level map: "xxxxxxxxxxxx";
  ```

Access Time Preservation

Two alternative statements are available for clients that do not allow *invisible file access*.

(With *invisible file access*: neither atime nor ctime is updated when files are backed up. Clients that support invisible file access: Solaris 2.1 or higher for SPARC, SunOS 4.1.3 for SPARC. This setting does not apply to these clients.)

For all other clients, you can decide whether each night's backup process itself either:

- preserves the original change time (ctime), but updates each file's access time (atime) to the time of the backup, or
- preserves the original access time (atime), but updates each file's change time (ctime) to the time of the backup

To preserve the original ctime, use
```
preserve file change time during backup;
```

To preserve the original atime, use
```
reset file access time after reading file;
```

Note: With `"reset file access time after reading file;"` incremental backups do not check the ctime for changes as they do with `"preserve file change times during backup;"`. The result: for applicable platforms, if you perform functions that change the inode, such as **chmod**, these changes are not recorded during incremental backups.

These two options are mutually exclusive. (If both are applied, a semantic warning is issued during backup processing and then the system's default operation — preserving ctime but updating atime — is performed.)

Maximum Files not Backed Up Before Forcing Full Backup

This statement directs catalog processing to schedule a full (level 0) backup (instead of an incremental) for this filesystem work item when it detects too many files within this work item that have never been backed up. (Applies to filesystem backups only, as database backups are always full.)

The concern is for files that do not get backed up during an incremental backup because the files were added in a manner that preserved the creation and access time of the file prior to the last backup of the work item. In this case, you want the next backup for that work item to be a level 0. Failure to do this causes catalog processing to take a long time and could result in a large number of files not being backed up, meaning they could not be recovered.

Also, see "When You Change a Work Item or a Filesystem" on page B-35 for times you should force a level 0 backup.

The format of the field is as follows:

```
maximum files not backed up before forcing full
backup: value;
```

For example:

```
maximum files not backed up before forcing full
backup: 10;
```

If **ebcatcomp** detects more than *value* files that have never been backed up for this filesystem work item, it schedules (command line scheduling) a level 0 for this work item. If *value* is set to zero, the default value, this feature is disabled.

Note: This statement can increase the number of level 0 backups performed.

This statement overrides a version of this directive that can be set in the server block, which applies to any filesystem work item. See "Catalog Threshold to Force Level 0 Backup" on page B-29.

Database Work Item Fields

Note: These "Database Work Items" pertain to the so-called "offline database backup" functionality supported prior to EDM 4.4.0 and are included here to support restores only. These work item specifications do not apply to the various database backup clients that are currently supported for backup.

Support for new backups using the "offline database backup" functionality are not supported as of EDM 4.4.0. Nor does EDM 4.4.0 support any reconfiguration of these database work items or configuration of new ones.

However, support for restore of backups taken using the "offline database backup" functionality continues in EDM 4.4.0, and to be able to restore such a backup depends on the continued presence of its particular database work item in your eb.cfg. Hence this section continues to be included in this revision of the *EDM Software Reference* manual.

Note: The work items that pertain to the various database clients are documented in the individual database client guides as appropriate; they are not included in this appendix.

A database work item defines a set of databases to be backed up. The following is an example.

```
work item: "chipmunk:master+:sybase", "chipmunk"

{

filespec:"DO_FILE_LIST /ext_ibm/sybase/master.dat";
database type:  "sybase";
database server:  "SYBCHIP";
database name:  "master model tempdb";
inclusion tag:  "chipmunk:SYBCHIP";
type: coordinated;
backup client initialization command: "-exec_as
  'sybase' -DBServer 'SYBCHIP' ebcv_shutdown
  sybase";
backup client cleanup command: "-exec_as 'sybase'
  DBServer 'SYBCHIP' ebcv_startup sybase";
use connection method: netware@514;
}
```

In the discussions that follow, the fields are covered in order as shown in the example above. The *work item* specifications are repeated as many times as necessary to configure backups for each client.

CAUTION: **These database work items must remain as they are for restores of their corresponding backups to work.**

Database Work Item Name

The database work item name that is generated during auto-discovery is of the format *database_name*[+]:*database_type*[:n]. If multiple databases were backed up in a work item, then *database_name* represents one of the database names (chosen arbitrarily), and a plus sign (+) is appended to the name. *database_type* is one of oracle, informix, or sybase.

Because a database work item name must be unique, an integer was appended if necessary to make it unique.

Client Name

This is name of the client to which the work item applies, as it appears in the authorized backup list. Always use the primary complete name for the client; not an alias.

```
work item: "doc1-all", "doc1"
```

Client Name

Filespec

The DO_FILE_LIST in this field specifies a file on the client that lists each file associated with the database that was found during database auto-discovery. This field has no character limit.

filespec:"DO_FILE_LIST /ext_ibm/sybase/master.dat"

Partition Spec

The DO_PART_LIST specifies a file on the client that lists each raw partition associated with the database that was found during database auto-discovery.

Database Type

This syntax specifies the database type:

database type: "*database_type*";

Database type is one of: informix, oracle, sybase, or none.

Database Server

This field sets the name of the database server for the work item. The syntax is:

database server: "*database_server_name*";

Database Name

This field contains one or more database names, separated by spaces.

Type

For a database work item, this field must be as follows:

type: coordinated;

Exclusion Tag

This field identifies any work item that should not be backed up concurrently with one or more other work items. For each client, the same tag value (1-63 characters) is specified for all those work items that are mutually exclusive, using this format:

```
exclusion tag: "tag";
```

For more information, see "Exclusion Tag" on page B-40.

Inclusion Tag

A single database might be backed up by two work items: one for the raw partition, and another for the files. The inclusion tag is used to coordinate back up of work items that refer to the same database. By default the inclusion tag value is the work item name. The syntax is:

inclusion tag: *"inclusion_tag"*;

Access Time Preservation

Two alternative statements are available for clients that do not allow *invisible file access*.

(With *invisible file access*; neither atime nor ctime is updated when files are backed up. Clients that support invisible file access: Solaris 2.1 or higher for SPARC, SunOS 4.1.3 for SPARC. This setting does not apply to these clients.)

For all other clients, you can decide whether each night's backup process itself either:

- preserves the original change time (ctime), but updates each file's access time (atime) to the time of the backup, or
- preserves the original access time (atime), but updates each file's change time (ctime) to the time of the backup

To preserve the original ctime, use:

```
preserve file change times during backup;
```

To preserve the original atime, use:

```
reset file access time after reading file;
```

These two options are mutually exclusive. (If both are applied, a semantic warning is issued during backup processing and then the system's default operation — preserving ctime but updating atime — is performed.

Priority

This field assigns a priority to the work item, forcing it to run either before or after one or more other work items. See "Priority" on page B-41 for more information.

Do Not Load Balance

Load balancing schedules backups as evenly as possible across all clients. The following statement excludes the work item from load balancing:

```
do not load balance;
```

For more information, see the discussion under File System Work Item Fields, "Do Not Load Balance" on page B-43.

Backup Client Initialization Command

This is the **ebcv_shutdown** command that is run on the database client before shutdown.

The command parameters are

```
-exec_as username \
-DBServer servername \
-DBA database_administrator_name \
-DBAIdent encrypted_dba_password \
database_type
```

Do not attempt to edit the encrypted password.

Initialization Timeout

This field specifies the amount of time to wait before the shutdown and backup are considered to have failed and the attempt is terminated. The syntax is:

```
backup client initialization timeout: n day n hour n minute n second;
```

Backup Client Cleanup Command

This is the **ebcv_startup** command that is run on the database client before restarting the database.

The command parameters are:

-exec_as *username* \
-DBServer *servername* \
-DBA *database_administrator_name* \
-DBAIdent *encrypted_dba_password* \
database_type

Do not attempt to edit the encrypted password.

Cleanup Timeout

Amount of time to wait before the restart is considered to have failed and is terminated. Note that this has no effect on the actual backup. The syntax is:

backup client cleanup timeout: *n* day *n* hour *n* minute *n* second;

Buffer Sizes

The buffer sizes for the server and clients can be tuned for better backup and restore performance. The buffer sizes you choose depend on your site's configuration. Factors that influence the optimum buffer size settings are: the number and type of library units and RAM in the backup server.

The syntax of the buffer size fields is:

backup client data buffer size: *n* MB;
backup server data buffer size: *n* MB;
recovery client data buffer size: *n* MB;
recovery server data buffer size: *n* MB;

Backup Start Time

This field specifies an additional qualification to the start time for nightly processing, which is specified in a crontab entry. The default crontab start of nightly processing is 11:00 p.m.

Level Map

This option overrides the backup level(s) that would normally occur for the work item, based on the template used.

For more information, see "Specify the completeness-code as any of the values below. This field applies for level 0-9 backups. The initial setting varies depending on the type of file that is being backed up (as noted in Table B-4). Generally it is a good idea to use the defaults." on page B-45.

PC Work Item Fields

These fields are used for all PC clients: NetWare, OS2, and Windows NT Backup Clients and for OpenVMS Backup Clients. In several cases the word netware is embedded in the code, but the field is used for all PC clients.

Note: For details, see the appropriate client manual.

The following is a sample Windows NT work item.

```
work item: "prince:/C/", "prince"
       {
        filespec: "DO_FS /C/";

        exclusion tag: "PRINCE_Harddisk0";

        use connection method: netware@3895;

        netware username: "supervisor";
```

Note: The "DO_FS" directive is not valid for NetWare and OS/2.

Work Item Name

This field defines the files to be backed up from a single client's resource. For example, the default syntax for Windows NT is:

workitem:"*servername:resource*","*servername*"

servername is the name of the PC server. *resource* is the PC file, directory, or volume to be backed up.

Connection Method

The connection method defines the TCP/IP port on which the client is listening for a command from the backup server. The syntax is:

use connection method: netware@*nnnn*;

where *nnnn* is the port number.

NetWare default	netware@1776
OS/2 default	netware@1492
Windows NT default	netware@3895
OpenVMS default	netware@3896
NT SQL Server	netware@5600

It must be the same as the port number that is set on the PC server. If you change one, you must change them both and rerun the configuration on the PC server.

Filespec

File specification describes the files to be backed up on the client.

Netware Username

The netware username entry determines what PC user privilege can execute backups and restores. For detailed information on how this works for individual clients, see the appropriate client manual.

Netware Encrypted Password

The netware encrypted password entry provides the required password to the PC user who can execute backups and restores. Do not attempt to edit it directly. Make any changes through the EDM Backup Configuration window.

Netware Client TSA

The netware client TSA entry assigns a target service agent to the PC server.

In NetWare, OS/2, and Windows NT the format is:
`servername.filesystem`

Netware Client Target

The netware client target defines a PC client as a target in need of service.

Exclusion Tag

The exclusion tag ensures single-threaded processing for the client. All work items with the same exclusion tag execute sequentially. Backup processes proceed one at a time.

Backup Trailsets

A backup media set (trailset) defines all of the trails that are used for one or more work groups, (including the type of media to which the backups should be written for each level). It defines the retention period for the backed up data and related catalogs and saveset records. It also defines the level at which you want to compress the catalogs (to save disk space).

When running **ebbackup**, you identify the schedule template to use. That template references to other types of configuration constructs:

- one or more work groups
- the trailset(s) that you want to use to back up those work groups

Each template requires one trailset (called the *primary trailset*), and can include a second (*alternate*) trailset. Generally the alternate trailset is used for off-site storage.

The backup trailset is similar to the following (for simplicity, only information for levels 0 and 9 are used):

```
backup trailset: "on site 1"
{
use trail: "fulls-tape", dlt for level 0,
using at most 8 clients concurrently;

duplicate media after backup on 1 copy,

appending to current media copy;
use trail: "incr-dlt", dlt for level 1-9,
using at most 8 clients concurrently;
use trail: "baselines", EO for level B1,
using at most 8 clients concurrently;
use level B1 for baseline backups;
keep backups of level 0 for: 1 year;
keep backups of level 9 for: 3 months;

keep duplicates of level 0 for: 1 year;

keep duplicates of level 9 for: 3 months;
keep backup catalogs of level 0 for: 1 year;
keep backup catalogs of level 9 for: 3 months;
keep saveset records of level 0 for: 1 year;
```

Note: Duplicate media information (as shown in lines 5, 6, 14, and 15 of the above example) appears only if media duplication is configured. If new media is requested for duplication, the line "using new media at each duplication" appears in place of line 6. If manual duplication is enabled, the line "manual activation of media duplication" appears after line 6.

Specify as many media sets (trailsets) as necessary to configure your site: one for each unique combination of backup levels, media, and length of time to keep the backup data, catalogs, and saveset records.

> **Note:** The fewer backup schedule templates, trailsets, and trails you use to configure your site, the fewer backup volumes EDM Backup needs to access during processing.

Define each field as described in the sections that follow.

Backup Trailset Name

Use this field to specify a 1-63 character name for the trailset (defaults to *primary* at installation). This name must be unique within the configuration file:

```
backup trailset: "trailset name"
```

> **Note:** The trailset name field does not end with a semicolon, but instead is followed by a brace-delimited block that defines its trails.

Changes to the trailset name take effect the next time **cbbackup** processing runs.

Use Trail

Use this field to specify the trail, including the name for the collection of media designated by the trail, the type of media that is used, the level of backups written to the trail, and the maximum number of clients that can be backed up concurrently to this trail (vs. the maximum number of clients that can be backed up concurrently for the entire server, and specified by using the global *maximum simultaneous clients* field in the server block).

You can have up to 12 trails per trailset — one for each backup level (0-9, B1, and B2). Each baseline level requires its own trail. Levels 0-9 can share trails. You generally want one or two trails for levels 0-9, plus the baseline trail(s).

This field can take two forms, and includes four fields: the trail name, media type, backup level, and number of concurrent clients. Each field is described following the syntax below:

```
use trail: "trailname", media-type for level n,
using any number of clients concurrently;
```

or

```
use trail: "trail name", media-type for level n,
using at most n clients concurrently;
```

Changes to the *use trail* portion of the field take effect the next time **ebbackup** processing runs. Changes to the *using ... clients concurrently* portion take effect immediately.

Trail Name

Specify the trail name as 1-11 characters. Trail names are used in volume management as part of the volume label name. For levels 0-9, the trailname and the media type are combined to form the label name. For baselines, the trail name is the entire label name.

Note: The more trails you use for each set of work groups that are backed up together, the more overhead you'll incur in terms of media management and tracking.

Media Type

Specify the type of media to which the trail is written.

Backup Level

Specify the level(s) of backup you want to write to this trail, using one of three formats:

- A single level (0 through 9, B1, or B2):
  `use trail: "trail_1", dlt for level 9, ...`
- A range of levels:
  `use trail: "trail_1", dlt for levels 0-9 ...`
- A compound (AND) statement that includes two levels, or ranges of levels:

```
use trail: "trail_1", dlt for levels 5 and
levels 8-9 ...
```

Only use level B1 or B2 if you purchased HSM, to indicate whether you want to maintain a level B1 or B2 backup. If you want to maintain two baseline levels, specify one trail as B1 and the other as B2.

> **Note:** Never write level B1 and B2 backups to the same trail.

If you identify more than one trail for a particular level, the last specification overrides the others. In the following example, EDM Backup uses the *incr-dlt* trail for level 9 and *mid-dlt* for level 5, even though levels 9 and 5 were included in the range of levels specified for the (first) *fulls-tape* trail:

```
use trail: "fulls-tape", dlt for levels 0-9,
using at most 8 clients concurrently;
use trail: "incr-dlt", dlt for level 9,
using at most 8 clients concurrently;
use trail: "mid-dlt", dlt for level 5,
using any number of clients concurrently;
```

Make sure to include every level of backup you want to run in at least one trail. For example, if you want to perform level-5 backups from the command line, you must define level 5 somewhere in the trailset. It is best always to include a range of levels (for example, 0 through 9) for each trail, even though some levels may never be used. (Automatically scheduled backups only use levels 0 and 9, for example.)

Maximum Number of Concurrent Clients

Specify the maximum number of work items (not clients) that can be backed up at the same time to this trail, in the range 1-24 (defaults to 8 at installation):

```
using at most n clients concurrently;
```

Alternatively, specify as many work items as possible (up to the number indicated by the server block's *maximum simultaneous clients* field):

```
using any number of clients concurrently;
```

The recommended setting is from 2 to 8:

- Use a lower number (for example, 2-3) if the trail is used mostly for full backups, or for incrementals that include a lot of file data (vs. extended inodes only)
- Use a higher number (up to 8) if the trail is used mostly for sparse incremental backups
- Choose a value in between if the trail is used for some combination of full backups and sparse incremental backups

This setting is used in conjunction with the system-wide (server block) *maximum simultaneous clients* field, to control the use of system resources. Regardless of the trail-level setting, there will never be more work items processed concurrently than are specified by the system-wide value.

For example, if you are using two trails that each specify four work items but your system-wide work-item limit is six, there can never be more than six work items writing to the two trails combined, at any one time. In this situation, EDM Backup:

- Starts backups on four work items using the first trail.
- Starts backups on two more work items using the second trail, bringing the total number of work items processed by the server to the limit.
- Starts a new work item each time an active work items finishes, keeping the total number of work items to six. In juggling the work items, EDM Backup does its best to allot equal resources to all scheduled work items.

This example is illustrated in Figure B-2 (where, for simplicity, assume one work item per client).

Figure B-2 Backup Processing

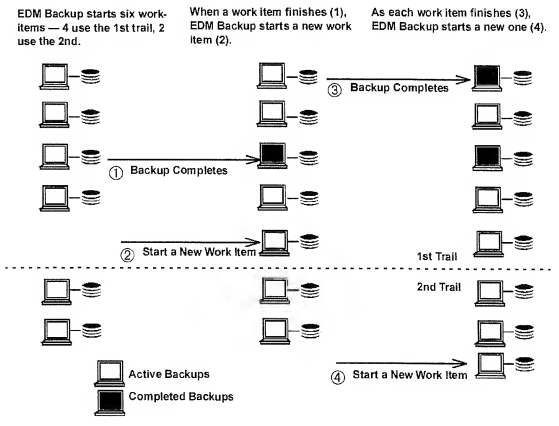

Use Level *B1* / *B2* For Baseline Backups

If you are using automatic scheduling, the scheduler writes a baseline backup to a B1 level trail. If an alternate trailset is configured with a B2 level trail, the autoscheduler writes an alternate night baseline backup to a B2 level trail.

Note: The entry `"use level B1 for baseline backups;"` in the eb.cfg file is ignored by the backup scheduler. The B2 level is always used for baselines on alternate night as long as the alternate trailset is configured with a B2 baseline trail.

Specify at most one baseline level per trailset. If you include this field, make sure to specify a *use trail* statement for the baseline backup.

Omit this field if you are not writing baseline backups to the trailset, or if you plan to request all baseline backups via custom or command-line scheduling. This field is ignored with custom and command-line scheduling.

Changes to this field take effect the next time EDM Backup schedules work items.

Keep Backups

Use this field to indicate how long to retain the backup data before reusing the backup media. Specify this field by using the following format, repeating it as many times as necessary to configure expirations for your site:

Note: If you identify more than one trail for a particular level, the last specification overrides the others.

```
keep backups of level n for: time-period;
```

where:

- *n* specifies the backup level (s) to which the setting applies (0-9), by using one of three formats as described for the *use trail* statement (a single level, range of levels, or compound statement that combines levels).

  Generally you define the same levels (or set of levels) here as are specified via *use trail* statements for the same trailset.

  Note: This specification does not apply for baseline backups.

- *time-period* defines the amount of time to retain the backups for this level, specified as an integer followed by an appropriate *units* value.

  Units

  second(s)

  day(s)

  week(s)

  month(s)

  year(s)

  forever (*do not include an integer*)

You can expire backups immediately:

```
keep backups of level 2 for: 1 second;
```

Or keep them indefinitely:

```
keep backups of level 0 for: forever;
```

You can combine the time units, as in the following example:

```
keep backups of level 9 for: 1 year; 6 months;
```

Always expire backup data before (or at the same time as) the corresponding saveset records, but after (or at the same time as) the corresponding catalogs. An exception occurs if you specify *keep backups forever*, in which case you can expire the saveset record without expiring the media. Make sure to cover every backup level that is used.

The recommended setting (and the initial value at installation) is 1 year for level-0 backups, and 3 months for levels 1-9:

```
keep backups of levels 0-9 for: 3 months;
keep backups of level 0 for: 1 year;
```

Note: If you identify more than one expiration for a particular level, the last specification overrides the others.

If you omit this field, it defaults to *forever*.

Changes to this field take effect immediately, but do not affect any currently running work items.

Keep Duplicates

Use this field to indicate how long to retain a duplicate of backup data before reusing the duplicate's media. Specify this field by using the following format, repeating it as many times as necessary to configure expirations for your site:

Note: If you identify more than one trail for a particular level, the last specification overrides the others.

```
keep duplicates of level n for: time-period;
```

where:

- *n* specifies the backup level(s) of the duplicate, to which the setting applies (0-9).
- *time-period* defines the amount of time to retain the duplicates for this level, specified as an integer followed by an appropriate *units* value.

Note: Refer to "Keep Backups" on page B-65 for more information about these fields.

You can expire duplicates immediately:

```
keep duplicates of level 2 for: 1 second;
```

Or keep them indefinitely:

```
keep duplicates of level 0 for: forever;
```

You can combine the time units, as in the following example:

```
keep duplicates of level 9 for: 1 year 6 months;
```

Note: You must expire duplicate data before (or at the same time as) the corresponding backup data. Refer to "Rejecting a Mount Request" on page 9-29 for more information.

Keep Backup Catalogs

Use this field to indicate how long to keep the backup catalogs created for each level 0-9 backup. Always expire the catalogs before, or at the same time as, the media (above).

The format of this field is:

```
keep backup catalogs of level n for: time-period;
```

Where *n* and *time-period* are specified as described above for the keep backups setting. Make sure to cover every backup level used.

The recommended setting (and the initial value at installation) is 1 year for level-0 backups, and 3 months for levels 1-9:

```
keep backup catalogs of levels 0-9 for: 3 months;
keep backup catalogs of level 0 for: 1 year;
```

Note: If you identify more than one expiration for a particular level, the last specification overrides the others.

You should expire the catalogs for incremental backups aggressively to free up disk space, because:

1. You rarely restore files from older incremental backups and
2. You can reconstruct the catalogs from the backups, if necessary

If you omit this field, it defaults to *forever*.

Note: See Chapter 10 "Magnetic Disk Concepts".

Changes to this field take effect immediately, but don't affect any currently running work items.

Keep Saveset Records

Use this field to indicate how long to keep the saveset records created for work items written to this trailset. Saveset records are necessary during a restore, to locate backup data and catalogs. Because of this, you can't generally expire saveset

records before either the backup data or the catalogs. An exception occurs if you specify *keep backups forever*, in which case you can expire the saveset record.

The format of this field is:

```
keep saveset records of level n for: time-period;
```

Where *n* and *time-period* are specified as described above for the keep backups setting, except that you can specify B1 and B2 as well as levels 0-9. Make sure to cover every backup level used.

The recommended setting (and the initial value at installation) is 1 year for level-0 and baseline backups, and 3 months for levels 1-9:

```
keep saveset records of levels 0-9 for: 3 months;
keep saveset records of level 0 for: 1 year;
keep saveset records of level B1 for: 1 year;
```

Note: If you identify more than one expiration for a particular level, the last specification overrides the others.

If you omit this field, it defaults to *forever*.

Changes to this field take effect immediately, but don't affect any currently running work items.

Backup Catalog Delta Level

Use this field to specify the backup level at which you want EDM Backup to consolidate catalogs (1-9). The recommended setting (and the initial value at installation) is 9:

```
backup catalog delta level: 9
```

If you omit this statement, the backup catalogs are consolidated at level 1.

When EDM Backup consolidates the catalogs, it turns all the catalogs for the level specified — as well as any numerically higher levels — into deltas, which only contain information that

differs from the more recent catalogs. If you specify a delta level of 5, EDM Backup compresses the catalogs for level 5-9 backups.

EDM Backup recreates the full catalog from a delta as necessary during a restore, by adding back the information that was originally compressed out. It uses as input all subsequent catalogs for the work item, up to (and including) the most recent catalog.

The latest catalog for any particular work item is always uncompressed, to provide fast access during restore processing. For the same reason, level-0 catalogs are always uncompressed.

Changes to this field take effect immediately, and could affect work items currently being backed up, but not any whose catalog is currently being processed. Since this field is used by catalog processing after backup processing, it should take effect on the next catalog to start being processed by **ebcatalogd** and **ebcatd**. This could be the currently running work items, since their catalogs have not been processed yet.

Backup Template Fields

A backup schedule template describes how to back up a group of clients, grouping together work group(s) and trailsets. You can have as many backup schedule templates as you want, but you must have at least one. The template specifies:

- The list of work group(s) to back up.
- The rotation period to use and date to start the first rotation.
- The trailset(s) on which to store the backup data for clients in the work group(s). Each backup template has a *primary* trailset, and may have an optional *alternate* trailset. Generally the alternate trailset is used for off-site storage.
- Whether to force all baseline backups to finish before any level 0-9 backups can start (available only if you've purchased HSM).

- Whether to recreate the baseline backup if necessary (available only if you've purchased HSM).
- Information about the log files and completion reports.
- Scheduling specifications.

Specify the template fields as described in the sections that follow.

Backup Template Name

Specify this field as the 1-63 character name for the template:
`backup template: "template name"`

> **Note:** The template name field does not end with a semicolon, but is followed by a brace-delimited block that defines its specifications.

The *template name* must be unique within the configuration file. You'll reference the template name when you invoke **ebbackup** to start a backup.

```
backup template: "sales-all"
{
work group list: "sales", "local";
begin trailset rotations on: Dec 13, 1993;
rotation period: 7 days;
primary trailset: "on site 1", use new volume on
each rotation;
alternate trailset: "off site 2", use new volume on
each rotation;
logging level: stats;
server logfile: "sales-all_template.log", 256k;
backup completion script: "mailok";
backup failure script: "mailerr";
do all baseline backups before normal backups;
recreate baseline if needed;
schedule:
{
 /* standard rotations; */
 /* full during weekends rotations; */
 weekday backup shift is 8 hours;
 weekend backup shift is 24 hours;
 level 9 on Monday, Tuesday, Wednesday, Thursday,
Friday;
 for "sales", level 0 on Saturday;
 for "local", level 0 on Sunday;
}
}
```

Note: If you remove a template from the configuration file, it remains in the schedule, as do any work items scheduled for that template. To remove a template completely, first remove it from the configuration file, then use **ebbackup** with the **-retire** option to delete it from the schedule.

Changes to the template name take effect the next time **ebbackup** processing runs.

Work Group List

Use the *work group list* to specify the work groups that are backed up together using this template. This is a shorthand method of identifying each work item individually. See "Work Group Fields" on page B-31 for information about defining work groups.

The format of this field is:

```
work group list: "work group", "work group",
...work group";
```

Changes to this field take effect the next time **ebbackup** processing runs.

Begin Trailset Rotations

This field specifies when EDM Backup should begin using the template (for a new template, it defaults to the date the template is added). It applies for all work items in the template, and for both automatic and custom scheduling (but not for command-line scheduling).

If you want to create a new template, but do not want to start using it yet, use this field with the data when the template is to be used. The format of this field is:

```
begin trailset rotations on: date;
```

Where *date* takes one of three formats:

Format	Example
dd-mmm-yy	13-dec-99
mm/dd/yy	12/13/99
mmm dd, yy	Dec 13, 99
mmm dd, yyyy	Dec 13, 1999

Leading zeros are not required. EDM Backup interprets the century using standard UNIX conventions.

- Years in the range 70-99 are in the 20th century (i.e., 19yy)
- Years in the range 0-37 are in the 21st century (i.e., 20yy)
- Years in the range 38-69 are illegal — EDM Backup can't create timestamps for these years because of a UNIX restriction

Changes to this field take effect the next time EDM Backup schedules work items.

Rotation Period

This field indicates how often each client should receive a level 0 backup (autoconfigured to every 14 days). For automatic scheduling, it's used by EDM Backup to compute the schedule of level 0 and level 9 backups, ensuring at least one level 0 backup during each rotation. For custom scheduling, it serves as the point of reference when defining the backup schedule (for example, run level 0 backups on the 1st day of each rotation, and run level 9 backups on days 4, 8, and 12).

The rotation period also indicates when EDM Backup should start a new media rotation for the template. By default, EDM Backup starts writing to a new set of media on day 1 of each new rotation. You can override this either in the statement that defines what trailsets to use (described under "Primary and Alternate Trailsets" on page B-75), or on the command line, when you run **ebbackup**.

The format of this field is shown below:

```
rotation period: n time-units;
```

Where:

- *n* is any integer greater than 1. The most common rotation periods are 7 days, 14 days, or 28 days. For automatic scheduling:
  - Use a 7-day rotation if restore time is important and you can afford the resources necessary to maintain a full backup every seven days.

- Use a 14-day rotation (the installation default) unless you have a good reason to use a different setting. The 14-day rotation conserves media and generally offers a good restore rate.
- Use a 28-day rotation to minimize the use of backup media, or for filesystems that change little over time (such as / and /usr). Restoring files will be slowest with this schedule, because there may be more incremental backups to read; however, the effect should be minimal if your data changes little over time.

- *time-units* defines the unit of time in which the rotation period is specified:

Units
day(s)
week(s)
month(s)
year(s)

For example, to automatically perform a full backup every 7 days on the clients named by the work group, enter:

```
rotation period: 7 days;
```

Changes to this field take effect the next time EDM Backup schedules work items.

Primary and Alternate Trailsets

A template must specify one trailset, and can specify two, known respectively as the *primary* and *alternate* trailsets.

Specify the trailset(s) using this format:

```
primary trailset: "trailset name",
use new volume on each rotation;
alternate trailset: "trailset name",
use new volume on each rotation;
```

Where:

- The *trailset name* identifies the trailset to which backups are written, and must specify a name that's defined already in the configuration file (set to "primary" at installation).
- The clause *use new volume on each rotation* directs EDM Backup to start a new volume(s) at the beginning of each rotation period. This separates backups from different rotation periods into distinct volume sets.

If you specify an alternate trailset, EDM Backup switches between the two trailsets each day, scheduling the primary trailset on odd-numbered days with respect to the template's rotation period, and the alternate trailset on even-numbered days. With automatic scheduling, this provides a second separate set of backups on alternating nights. If the rotation period specifies an even number of days, the two sets of backups are nearly the same, differing only by one day's worth of data.

With custom scheduling it can do the same, but it's up to you to schedule identical backups on alternating days. If you specify an alternate trailset with custom scheduling, make sure to define two complete (and exact) schedules, running the same level on each of two consecutive days. By scheduling the following levels to run on the days shown, you'll create two identical trailsets that differ by one day's worth of data.

Day	Level	Trailset Used
Monday	0	Primary
Tuesday	0	Alternate
Wednesday	9	Primary
Thursday	9	Alternate
Saturday	9	Primary
Sunday	9	Alternate

You can spread out the backups more, as long as you specify one of the two identical backups on an even day with respect to the rotation cycle, and one on an odd day.

Changes to this field take effect the next time **ebbackup** processing runs.

Logging Level

Use this field to specify the level of logging messages written to the file specified via the server log file field when this template is backed up.

Specify the logging level using this format:

```
logging level: logging level;
```

Where *logging level* is specified as follows (defaults to *stats* at installation).

Code	Description of Messages Logged
none	No messages.
errors	Logs errors, including device and communication errors.
stats	Logs statistics as well as the information for the *errors* level. The statistics include the amount of data saved, the time it took to save the data, when the backup started and finished, and so on.
debug	Logs debugging information as well as the information for the *stats* level. Only use this level at the direction of customer support.
per file	Logs several lines, including *debug*-level information, for each file that's backed up. Only use this level at the direction of customer support. It slows backup throughput and can result in a huge log file.

Changes to this field take effect immediately, but don't affect any currently running work items.

Server Log File

Use this field to specify the name and maximum size of the template's log file. This file is stored on the server as /usr/epoch/EB/log/*file-name* and records information as directed by the *logging level* field. It provides an audit trail of all backup activities for the template.

The format is:

```
server logfile: "log file name", file-size;
```

Where:

- *log file name* identifies the log file on the server (stored in the /usr/epoch/EB/log directory, set initially to default_template.log, where "default" is the template name). You can create a different log file for each template or you can share a single log file across all templates. If you specify a separate log file for each template, use a name that identifies the template.
- *file-size* indicates how large the file can get before the oldest data is expired. When the file reaches the specified size, EDM Backup locates the oldest data in the file and expires 10 percent of that data to free up space for new information. Specify this field as a number of bytes followed by a unit code (e.g., 500KB, set to 256K at installation). Specify the unit code as follows:

Unit Code	Measure
k, K, kb, or KB	kilobytes
m, M, mb, or MB	megabytes
g, G, gb, or GB	gigabytes

To allow the file to grow until it's as large as permitted by physical storage space, specify *no limit* (and no other options). By not setting a limit on the file, it becomes a permanent audit trail of backup operations for the template.

```
server logfile: "default_template.log", no limit;
```

Note: If you don't limit the file size, it may become too large to manage. Also, in HSM systems it won't be staged.

This example logs messages for the *sales-all* template, and has a maximum file size of 256KB:

```
server logfile: "sales-all_template.log", 256KB;
```

If you omit this setting, no log file is maintained for the template.

Changes to this field take effect immediately, but don't affect any currently-running work items.

Backup Completion Script

EDM Backup generates completion reports describing each successful backup. There is one completion report per template. EDM Backup forwards these reports to the script file indicated using this parameter. The script file, in turn, dispatches the reports according to how the script is defined.

As installed, EDM Backup directs completion reports to the supplied "mailok" script on the server (/usr/epoch/EB/config/mailok). Depending on how that script is defined, it sends the reports to specific administrators (generally those defined in the list of backup administrator usernames), to the log file, or to both.

You can change the script by entering a new script name in the **Advanced Options** popup window in the **Schedule** tab of the EDM Backup Configuration window.

If you specify a relative pathname, the file is assumed to be in /usr/epoch/EB/config. If you specify an absolute pathname, the file can be located anywhere.

For a sample report, see "Backup Completion Reports" on page 16-29.

If you this setting is blank, no completion reports are generated.

Changes to this field take effect immediately.

Backup Failure Script

EDM Backup creates backup failure reports for each work item whose backup fails. It forwards these reports to the script file indicated using this parameter. The script file, in turn, dispatches the reports according to how the script is defined.

By default, EDM Backup directs failure reports to the EDM-supplied "mailerr" script on the server (/usr/epoch/EB/config/mailerr). Depending on how that script is defined, it mails the reports to specific administrators (generally those defined in the list of backup administrator usernames), to the log file, or to both.

You can change the script by entering a new script name in the **Advanced Options** popup window in the **Schedule** tab of the EDM Backup Configuration window.

If you specify a relative pathname, the file is assumed to be in /usr/epoch/EB/config. If you specify an absolute pathname, the file can be located anywhere.

For a sample report, see "Backup Failure Reports" on page 16-31.

If this setting is blank, no failure reports are generated.

Note: You can leave this field blank to reduce the amount of EDM Backup mail. The backup completion script (*mailok*) reports include all types of activity, including failures. You can refer there for information about failures.

Changes to this field take effect immediately.

Do All Baseline Backups Before Normal Backups

Use this statement if you've purchased HSM, and want to ensure that *all* baseline backups finish for the template before the level 0-9 backups start for *any* work items backed up by the template:

```
do all baseline backups before normal backups;
```

This feature is useful when baseline backups consume most of the server's resources, which can happen if significant amounts of magnetic disk space are required to hold temporary files. It can also alleviate disk thrashing. On the downside, it can slow overall throughput.

With this statement omitted (the default at installation), EDM Backup finishes the baseline backup *for each work item* before starting the level 0-9 backup for that same work item.

Note: This specification applies for custom and automatic scheduling, but is not used for command-line scheduling.

Changes to this field take effect the next time **ebbackup** processing runs.

Recreate Baseline if Needed

Include this statement if you've purchased HSM, and you want to automatically re-baseline any file for which the existing baseline copy was substituted for the staged copy during a restore:

```
recreate baseline if needed;
```

If you don't want to re-baseline files automatically, use the following statement instead (or don't include either statement). This is the default setting at installation:

```
do not recreate baseline if needed;
```

If you choose to re-baseline files automatically, **ebbackup** checks each file when it runs a baseline backup, to see if the staging id in the file's extended inode is the same as the baseline id. These two ids are only the same if the baseline copy of the file was substituted for a missing (or damaged) staged version. If the ids are the same, **ebbackup** automatically generates a new baseline copy of the file, then updates the baseline id so that it points to the new copy.

Changes to this field take effect the next time **ebbackup** processing runs.

Schedule

The schedule tells EDM Backup when to run backups for the template, and how much time to commit to backup processing each day. For custom scheduling, it defines when to run each backup level (and optionally, when to run each level *for a specific work group*).

Some scheduling options apply for automatic scheduling and some for custom scheduling, as follows:

Table B-6 **Scheduling Fields**

Option	Applicable for:	
	Automatic	**Custom**
Standard rotations	yes	no
Full during weekends rotations	yes	no
Weekday backup shift	yes	no
Weekend backup shift	yes	no
[for *work-group-name*] level *n* on days ...	no	yes

The schedule block looks like this:

```
schedule:
{
/* standard rotations; */
/* full during weekends rotations; */
weekday backup shift is 8 hours;
weekend backup shift is 24 hours;
level 9 on Monday, Tuesday, Wednesday, Thursday,
Friday;
for "sales", level 0 on Saturday;
```

```
for "local", level 0 on Sunday;
}
```

Note: A level map takes precedence over any scheduling defined using this block. For example, if you'd normally create one level 0 backup and the rest level 9 backups during the rotation (as for automatic scheduling), but a work item specifies a level map of "B00000000000", each level 0-9 backup for that work item will run as a level 0 backup.

Specify each field as described below. In general, try to schedule backups when system use is low.

Standard vs. Full-During-Weekends Rotations

Most sites let EDM Backup schedule backups automatically, based on the settings in the configuration file. If this is the case at your site, choose one of these two scheduling directives to:

- Turn on automatic scheduling (selecting either directive does this).
- Specify how to perform backups. Indicate the statement you want by placing comment markers (/* ... */) around the other statement:

```
standard rotations;
/* full during weekends rotations; */
```

Specification	Description
Standard rotations	Schedules backups so that some portion of the clients receive a full backup each day
Full during weekends rotations	Schedules backups so that all clients that require a full backup receive that backup on Saturday or Sunday, when possible, with any incrementals running on weekdays

Changes to this field take effect the next time EDM Backup schedules work items.

Specifying Backup Shifts

If you're using automatic scheduling, use these two fields to indicate how much time you want to allow EDM Backup to run each day. This serves as a guideline, or goal, when EDM Backup schedules the work items for backup. Specify the time for each weekday (Monday-Friday), then for each weekend day (Saturday and Sunday):

```
weekday backup shift is hh hours mm minutes;
weekend backup shift is hh hours mm minutes;
```

Specify *hh* as a number of hours (an integer in the range 1-24), and *mm* as a number of minutes (1-60; include only if necessary):

```
weekday backup shift is 8 hours 30 minutes;
weekend backup shift is 24 hours;
```

> **Note:** The backup shift specifications are not binding, but serve only as a target for the amount of time committed to backup processing each day.

Changes to this field take effect the next time **ebbackup** processing runs.

Scheduling Custom Backups (Level n on Days...)

If you're scheduling custom backups, use this statement to:

- Turn on custom scheduling (happens automatically when you include one or more custom-scheduling statements and you omit both statements used to turn on automatic scheduling)
- Define when you want to run the backups, repeating this statement as necessary to configure your site:

```
[ for "work-group-name"] level n on [ day(s)]
days;
```

You can specify each work group individually, or you can combine the work groups for the template into a single statement. You might want to combine the work groups for one or more level(s), but separate them for another level(s):

```
level 9 on Monday, Tuesday, Wednesday, Thursday,
Friday;
```

```
for "sales", level 0 on Saturday;
for "local", level 0 on Sunday;
```

Specify each syntax field as follows:

- The *for work-group-name* clause applies if the statement is for a specific work group, and specifies the name of the work group exactly as it appears in the template's work-group list (defaults to all work groups in that list if you omit this clause). If two statements apply for the same work group, EDM Backup uses the most restrictive specification (the one that's specific to the work group). You might use one statement to schedule all work groups, then add statements for each exception to the "rule":
  ```
  level 9 on days 1-14;
  level 0 on days 7, 14;
  for "sales", level 0 on day 1;
  ```

- The level n clause identifies the backup level you want to run on the day(s) specified (0-9, B1, or B2). If more than one level is specified for the same day, EDM Backup performs the lowest level, numerically (e.g., level 0 instead of level 5).

  If the template specifies an alternate trailset, make sure to schedule two level-0 backups for each work group — one on an even-numbered day and one on an odd-numbered day, with respect to the rotation cycle. (Remember that the two trailsets are backed up on alternating days.)

  Note: Specify baseline backups explicitly. EDM Backup does not run baselines automatically when you use custom scheduling.

- *days* indicates the days on which you want to run backups for this level (and optionally, for this work group).

  If you don't request backups for a particular day, EDM Backup won't perform backups on that day (for this template).

Note: If you schedule custom backups, don't specify either of the directives for automatic scheduling (described under "Specifying Backup Shifts" on page B-84). This causes an error.

Specify the days as follows:

Days Specification	Description
a single integer	To represent that day in the rotation period (1 for the first day, 2 for the second day, and so forth up to the number of days specified in the rotation period for the template).
a range of integers	To represent a range of days in the rotation period *only for a rotation period this is NOT a multiple of* 7 *days long* (e.g., 1-6 for days 1, 2, 3, 4, 5, and 6)
name of day(s)	To represent a specific day of the week *(that is, for a rotation period of* 7 *days or a multiple of* 7 *days long).* Spell out the full name of the day (Saturday, Sunday, etc.). You *cannot* specify a range of names — each day must be listed separately. You can describe a specific occurrence of a day within the rotation period, by prefacing the day with a number (2nd Saturday, 3rd Monday, etc.).

Changes to this field take effect the next time EDM Backup schedules work items.

Startup Parameters

Use the startup parameters block to specify when to perform the first (level 0) backup for a new client. It's important to perform the first level 0 backup as soon as possible, because no incrementals run until the first level-0 backup is performed.

This block applies for automatic scheduling only, and is ignored with custom scheduling. It looks like this:

```
startup parameters:
{
perform initial full backup on scheduled day;
```

```
/* perform initial full backup as soon as possible;
*/
;
```

Specify one of the following in the EDM Backup Configuration window:

- Specify *perform initial full backup on scheduled day* to perform the initial full backup on some portion of the newly-installed clients each day during the first rotation period. This setting distributes the work load as evenly as possible across the rotation period, and is the recommended (and initially configured) approach. By the end of the first rotation, every client will have a level-0 backup.
- Specify *perform initial full backup as soon as possible* to run the initial full backup on all newly-installed clients during the first backup run after those clients are installed. When you use this option, EDM Backup performs initial full backups on every new client, one after the other. The impact on the system is dictated by the total amount of client data that needs to be backed up.

You may want to specify *perform initial full backup on scheduled day* when you first configure your system, then change it to *perform initial full backup as soon as possible* after the initial backups are taken (for use as new clients are added).

With either method, the client gets backed up according to its normal rotation schedule after the initial level-0 backup is completed.

> **Note:** This specification only applies for the first backup that's run after a client is installed. It has no effect on subsequent processing.

Changes to the startup parameters block take effect the next time backup scheduling occurs.